



Jason S Wong

<http://usa.redirectme.net>

Sr. DBA

IT Applications Manager

DBA Developer

Programmer

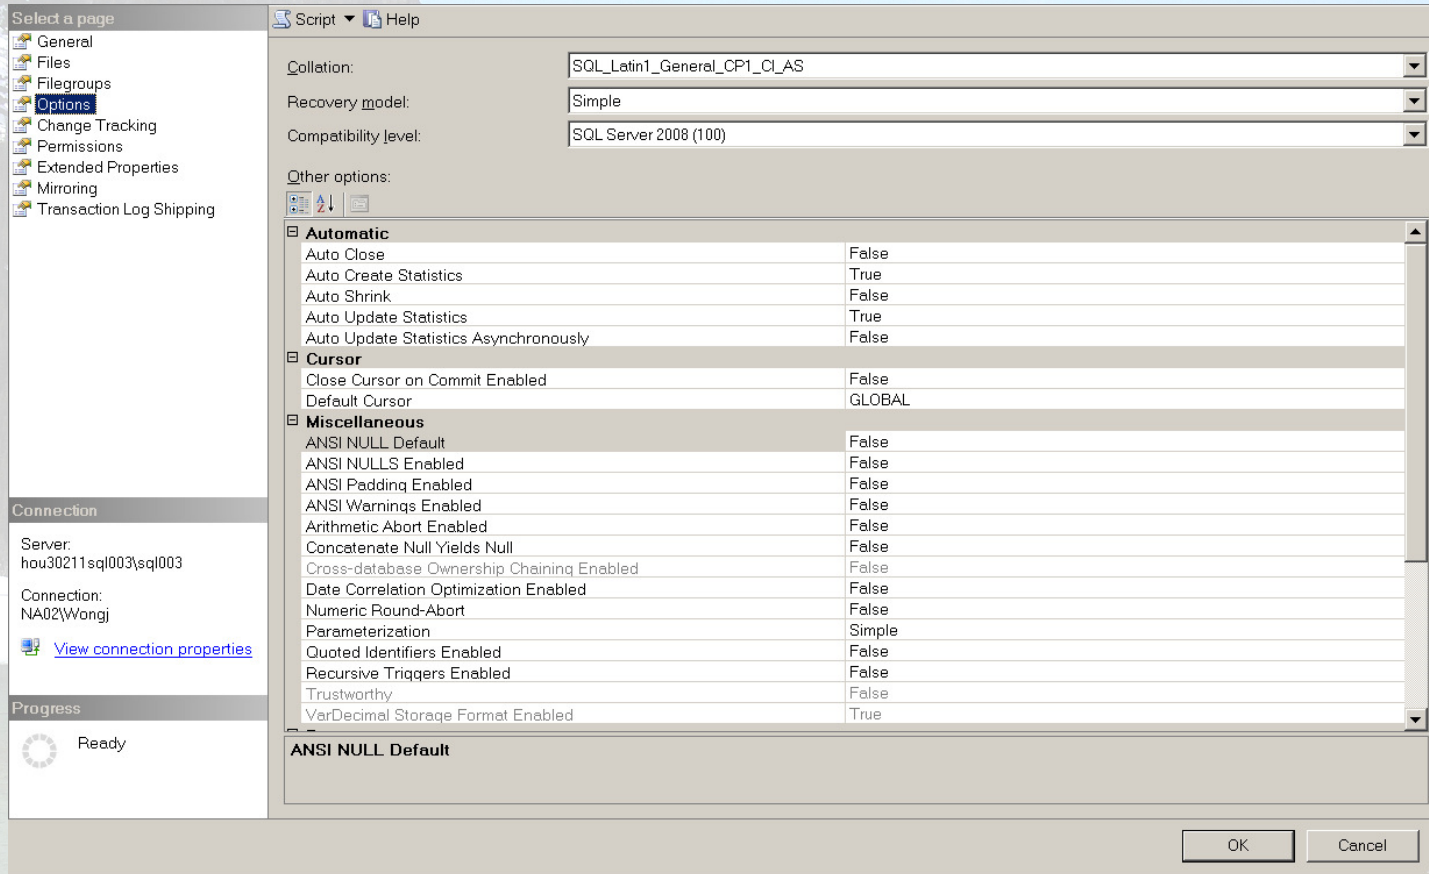
M.S. Rice '88, MBA U.H. '94 (MIS)

Make your defaults

Top 10 SQL Server defaults that DBAs need to evaluate and modify for their environment.

- Recovery model
- Min - Max memory
- Initial size, auto-growth
- Default Isolation Level Read Committed
- (Index) FillFactor
- Primary Key Clustered Index
- Uniqueidentifier
- Collation
- Data type - nvarchar(255), the "N"
- Allow Nulls

1. Recovery model




Full recovery mode without backup job;

QA/DEV refresh reset recovery model; bad queries fill the log on production host; no more disk space

Truncate_Log.sql - next

Back by date-time naming file script - next



```
EXEC sp_dboption 'db', 'trunc. log on chkpt. ', 'TRUE';
```

```
Checkpoint
```

```
EXEC sp_dboption 'db', 'trunc. log on chkpt. ', 'FALSE';
```

```
DBCC SHRINKFILE (db, 2);
```

```
GO
```

```
-----discontinued functionality in 2008-----
```

```
BACKUP LOG db_log WITH NO_LOG;
```

```
GO
```

```
SELECT log_reuse_wait_desc FROM sys.databases WHERE NAME = 'db_log';
```

```
GO
```

```
DBCC SHRINKFILE (db_log, 2);
```

```
GO
```

```
DO NOT DETACH DATABASE, DELETE LOG FILE, REBOOT
```




```
--Full backup using datetime file name
```

```
DECLARE @devname varchar(256)
```

```
DECLARE @DirLocationPre varchar(256)
```

```
DECLARE @DatabaseName varchar(256)
```

```
SET @DirLocationPre = 'U:\backup\DB1\DB1FullBackup'
```

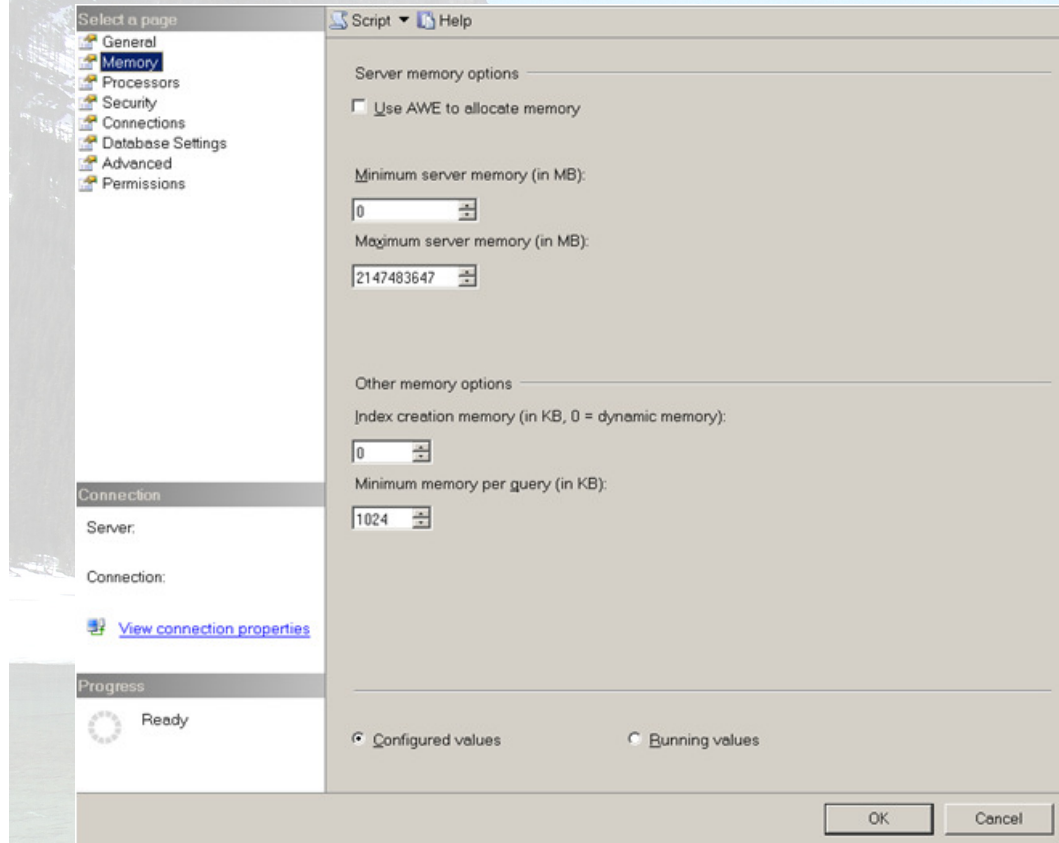
```
SET @DatabaseName = 'DB1'
```

```
SELECT @devname = @DirLocationPre + REPLACE(REPLACE(REPLACE(CONVERT(Varchar(40),  
GETDATE(), 120), '-', ''), ':', ''), ' ', '') + '.bak';
```

```
BACKUP DATABASE @DatabaseName to DISK = @devname
```

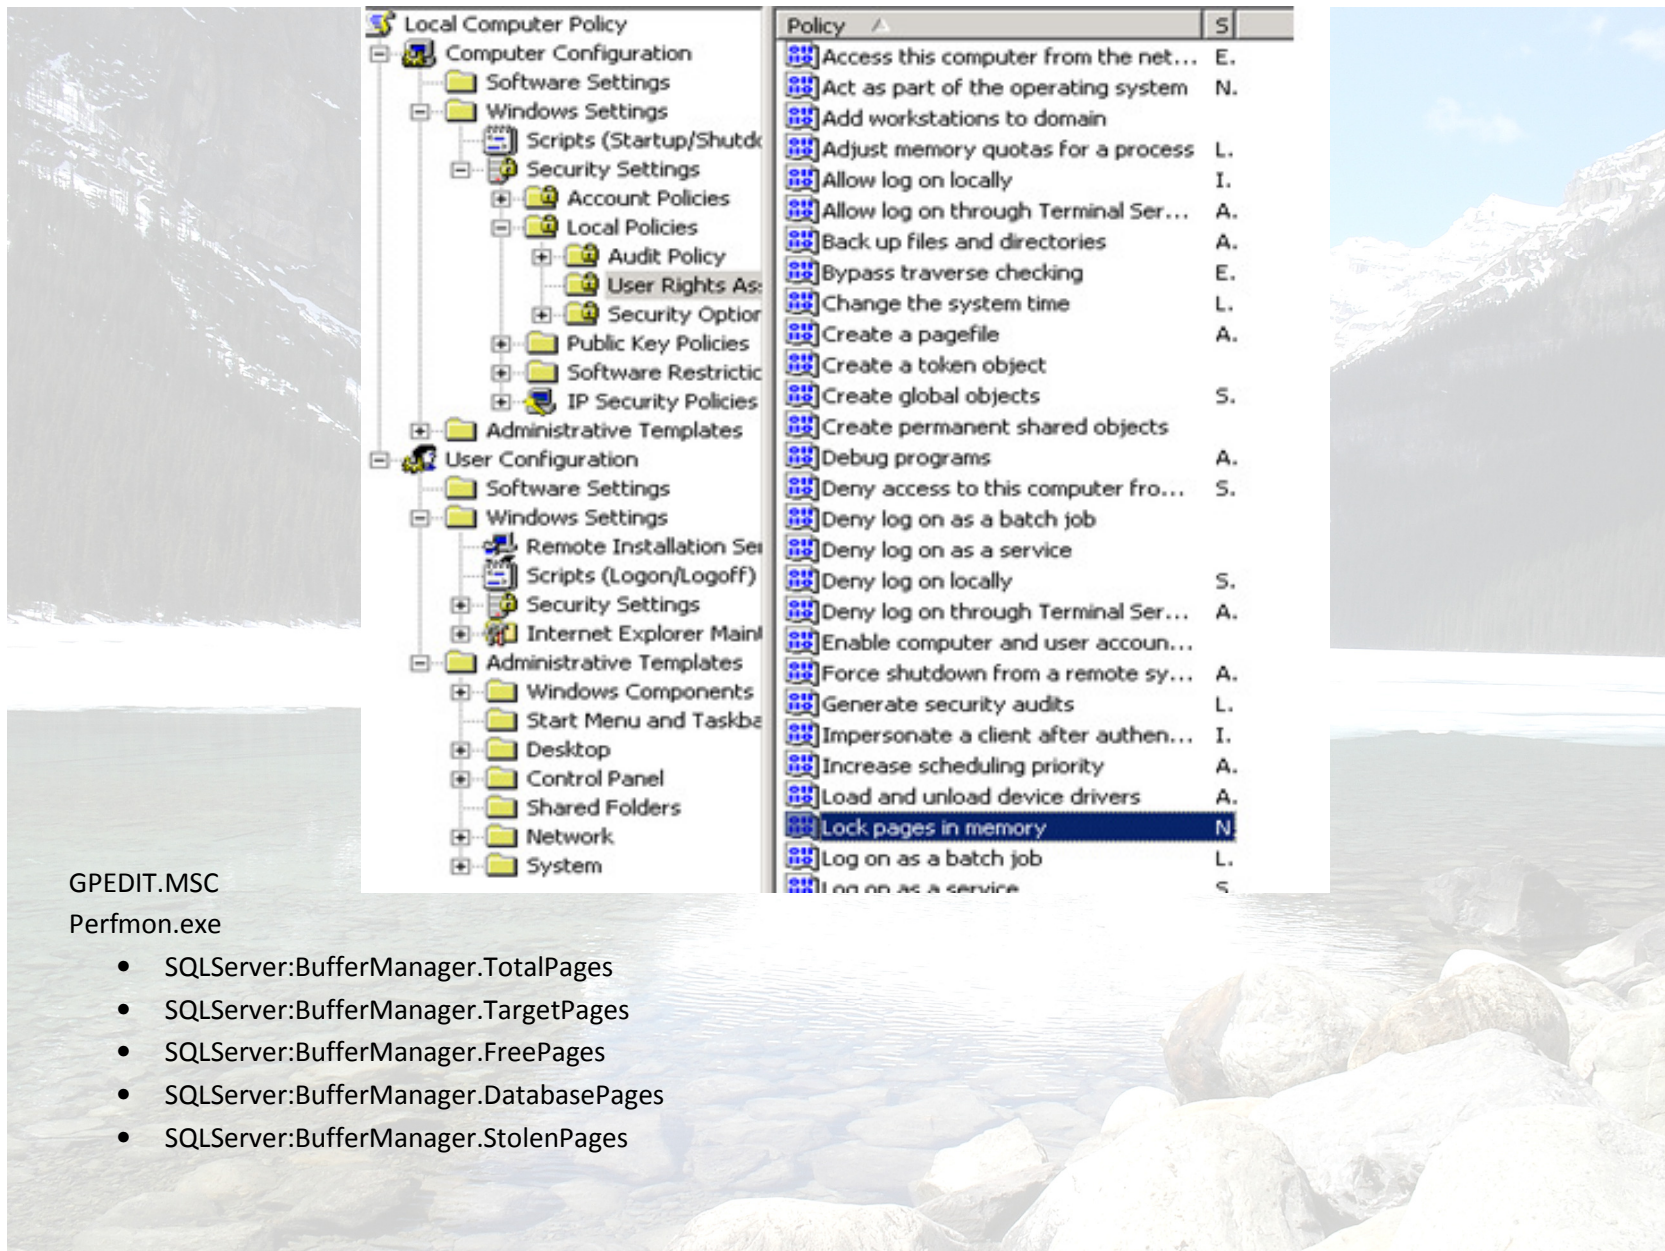
```
GO
```

2. Min – Max memory



Memory Usage:

- Buffer cache, procedure cache, dirty pages
- SPIDMemoryUsage.sql -- check paging and cache fault counters
- StalledDatabaseIO.sql -- check physical disk counters
- Parameterized SP vs ad-hoc query.
- Memory speed vs Disk IO speed



GPEDIT.MSC

Perfmon.exe

- SQLServer:BufferManager.TotalPages
- SQLServer:BufferManager.TargetPages
- SQLServer:BufferManager.FreePages
- SQLServer:BufferManager.DatabasePages
- SQLServer:BufferManager.StolenPages

3. Initial size, auto-growth

Database Properties - model

Select a page

- General
- Files
- Filegroups
- Options
- Change Tracking
- Permissions
- Extended Properties

Script Help

Database name: model

Owner: sa

Use full-text indexing

Database files:

Logical Na...	File Ty...	Filegroup	Initial Size (...)	Autogrowth	Path
modeldev	Rows ...	PRIMARY	2	By 1 MB, unrestricted growth	C:\Program Files\Microsoft SQL Server\MSSQL10.SQL...
modellog	Log	Not Applic...	1	By 10 percent, unrestricted gro...	C:\Program Files\Microsoft SQL Server\MSSQL10.SQL...

Connection

Server: P8661VF1\SQL2K8

Connection: NA02\Wongj

[View connection properties](#)

Progress

Ready

Add Remove

OK Cancel

Define realistic size and growth rate, 1K per row, 1 million records 1 GB

DiskSpaceReview_SP.sql

Default location of tempdb, system DB

RAID 1, 5, 10, number of disks and spindle speed

4. Default Isolation Level Read Committed

SELECT Shared Lock

- Table hint: WITH (NOLOCK) = Read Uncommitted
- ALTER DATABASE AdventureWorks SET ALLOW_SNAPSHOT_ISOLATION on
- ALTER DATABASE AdventureWorks SET READ_COMMITTED_SNAPSHOT on
- Data-warehouse reporting model vs transaction model
- Schedule jobs at quiet time

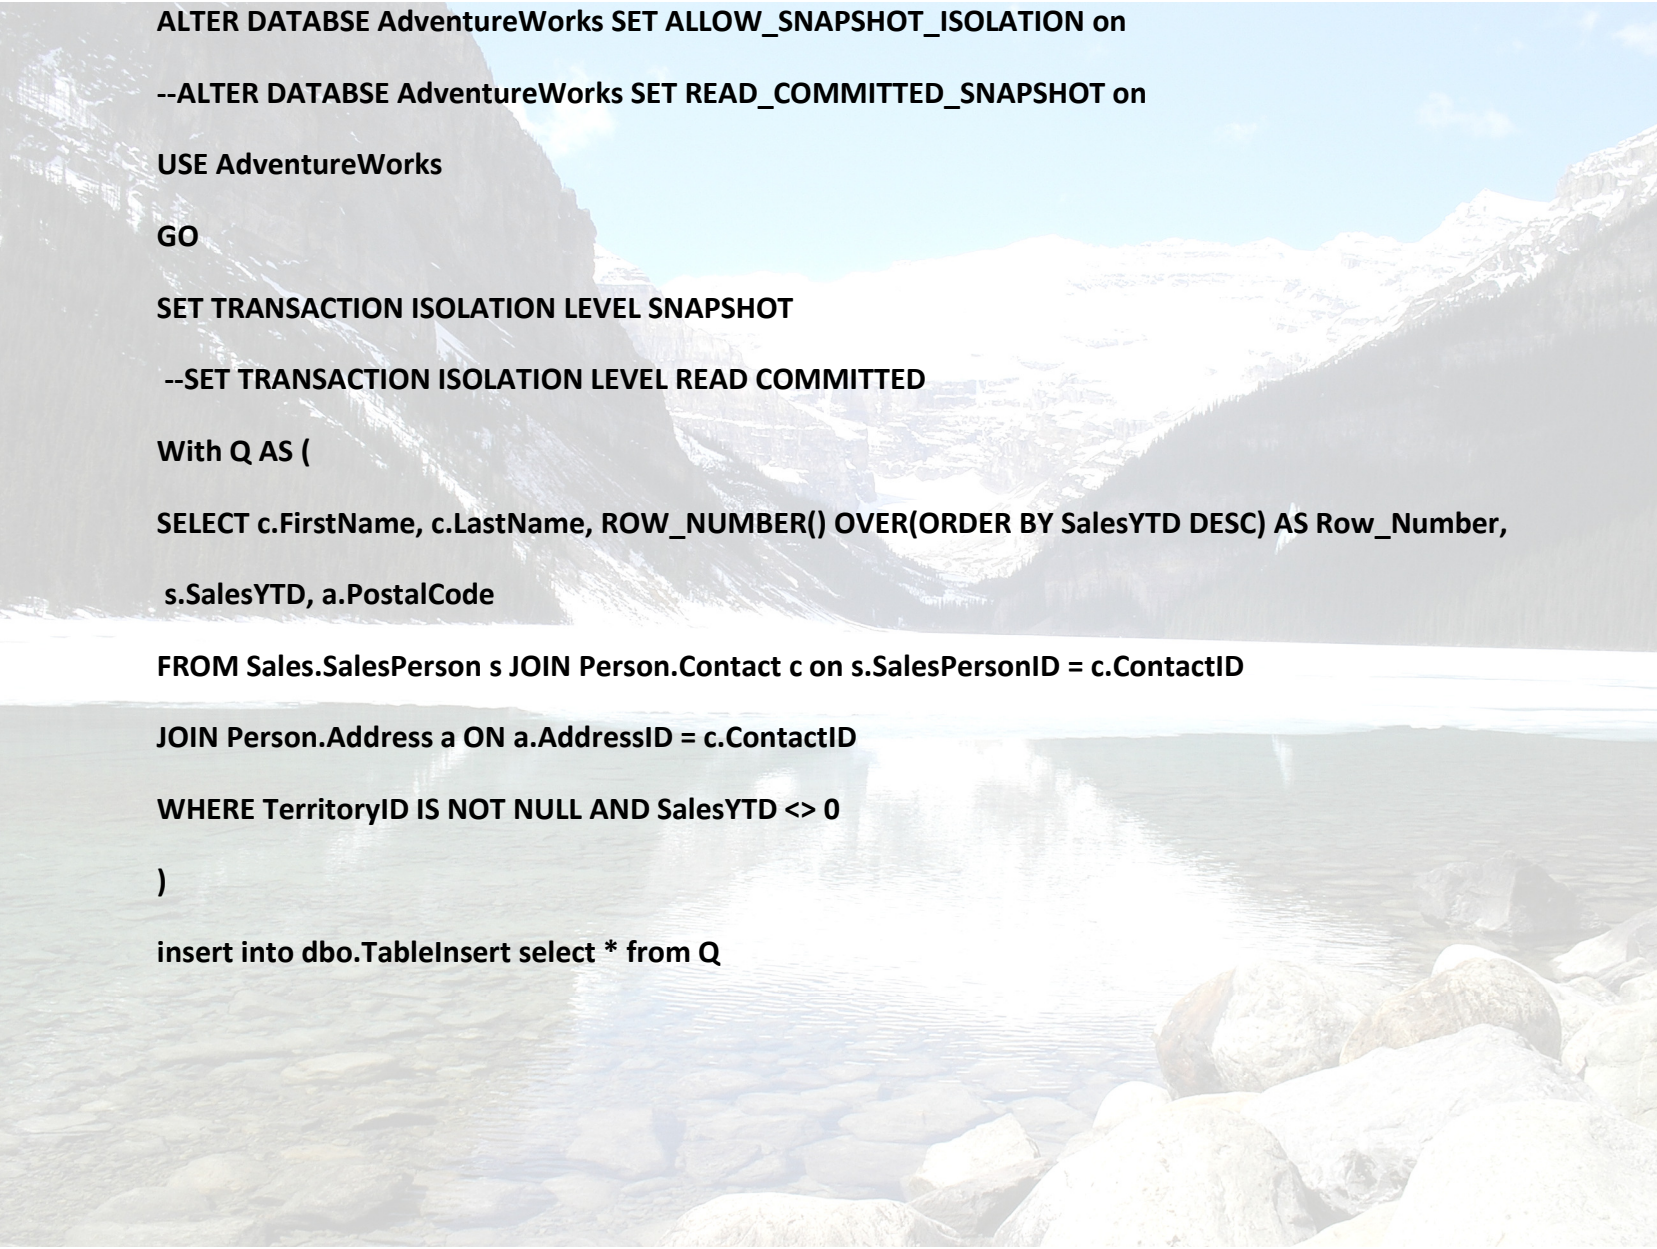
BOL – “Enabling Row Versioning-based Isolation Levels”

Row versioning without exclusive lock

SnapshotIsolation.sql - next

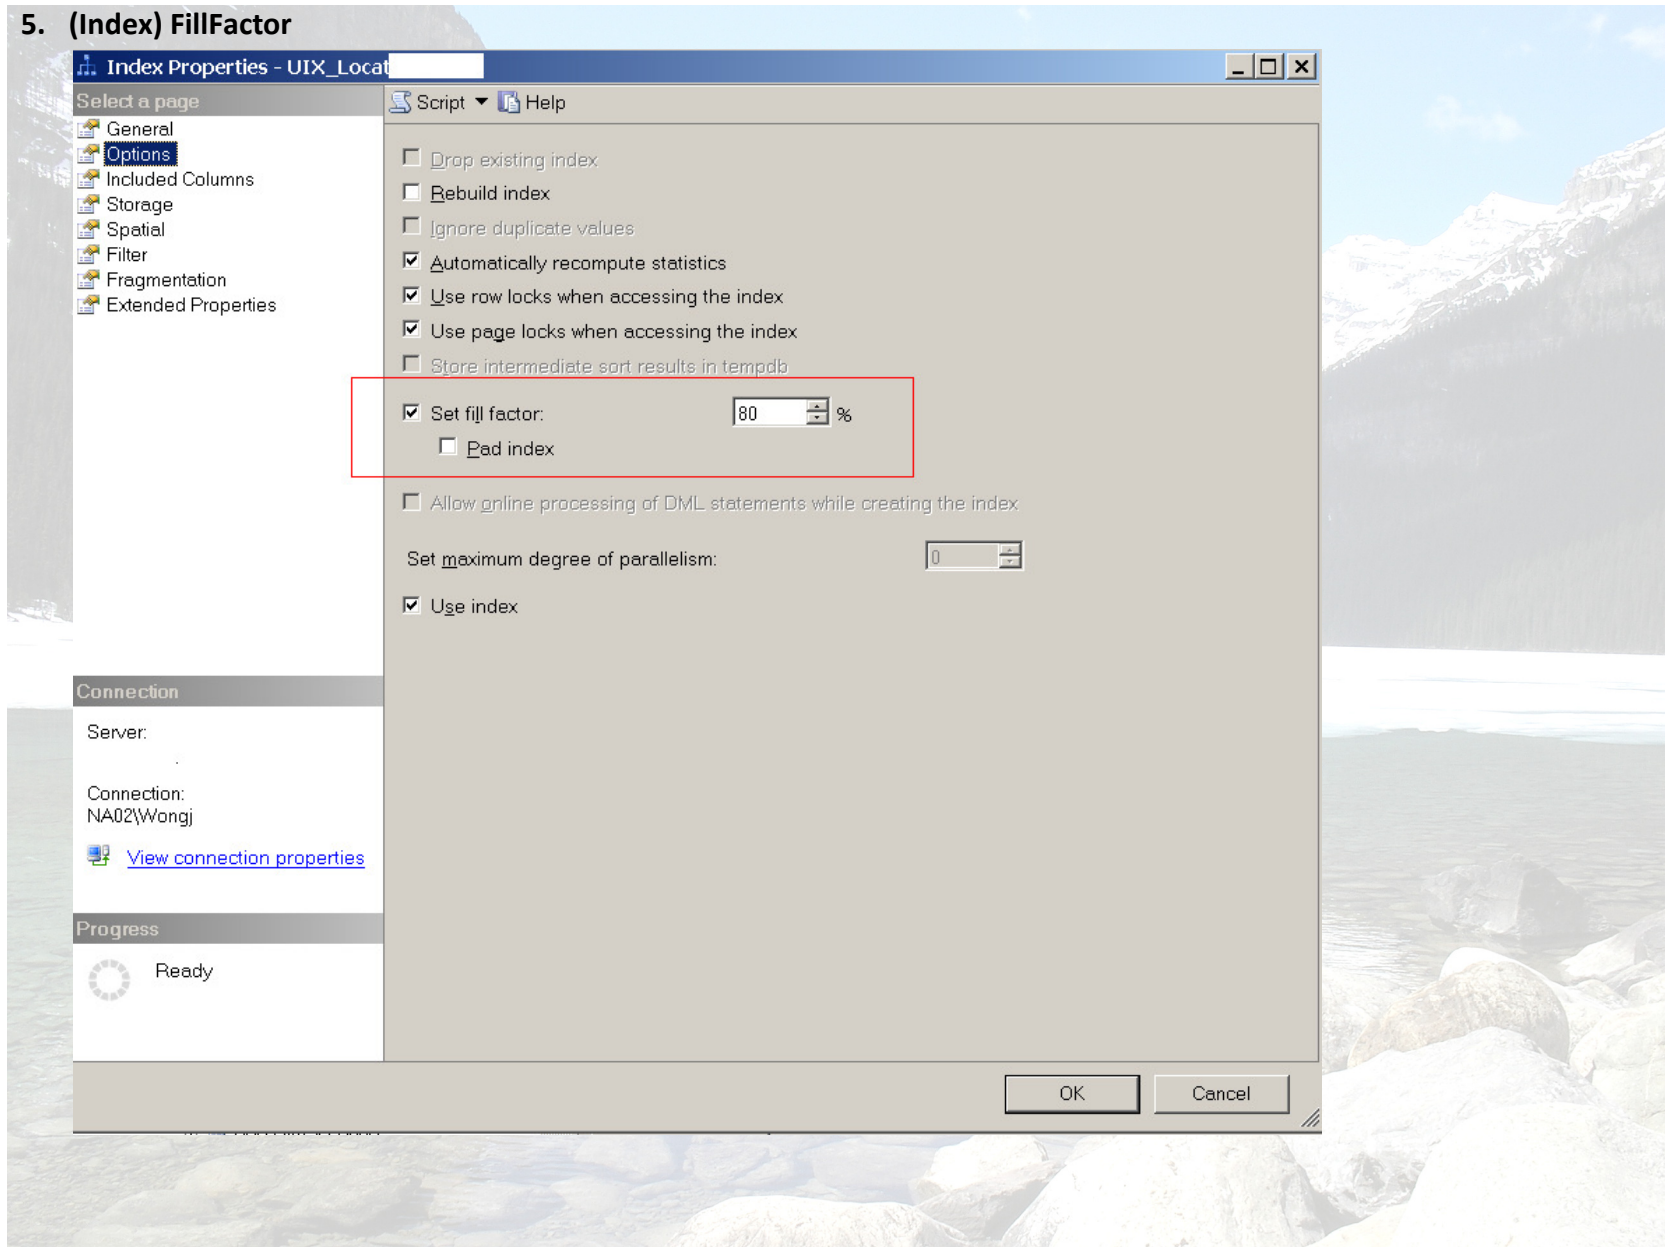
Concurrency side effects enabled by isolation levels

Isolation level	Dirty read	Nonrepeatable read	Phantom
Read uncommitted	Yes	Yes	Yes
Read committed	No	Yes	Yes
Repeatable read	No	No	Yes
Snapshot	No	No	No
Serializable	No	No	No



```
ALTER DATABASE AdventureWorks SET ALLOW_SNAPSHOT_ISOLATION on
--ALTER DATABASE AdventureWorks SET READ_COMMITTED_SNAPSHOT on
USE AdventureWorks
GO
SET TRANSACTION ISOLATION LEVEL SNAPSHOT
--SET TRANSACTION ISOLATION LEVEL READ COMMITTED
With Q AS (
SELECT c.FirstName, c.LastName, ROW_NUMBER() OVER(ORDER BY SalesYTD DESC) AS Row_Number,
s.SalesYTD, a.PostalCode
FROM Sales.SalesPerson s JOIN Person.Contact c on s.SalesPersonID = c.ContactID
JOIN Person.Address a ON a.AddressID = c.ContactID
WHERE TerritoryID IS NOT NULL AND SalesYTD <> 0
)
insert into dbo.TableInsert select * from Q
```


5. (Index) FillFactor




```
ALTER INDEX { index_name | ALL }
  ON <object>
  { REBUILD
    [ [ WITH ( <rebuild_index_option> [ ,...n ] ) ]
      | [ PARTITION = partition_number
          [ WITH ( <single_partition_rebuild_index_option>
                [ ,...n ] )
            ]
        ]
    ]
  | DISABLE
  | REORGANIZE
    [ PARTITION = partition_number ]
    [ WITH ( LOB_COMPACTION = { ON | OFF } ) ]
  | SET ( <set_index_option> [ ,...n ] )
  }
[ ; ]
```

```
<object> ::=
{
  [ database_name. [ schema_name ] . | schema_name. ]
  table_or_view_name
}
```

```
<rebuild_index_option > ::=
{
  PAD INDEX = { ON | OFF }
  | FILLFACTOR = fillfactor
  | SORT_IN_TEMPDB = { ON | OFF }
  | IGNORE_DUP_KEY = { ON | OFF }
  | STATISTICS_NORECOMPUTE = { ON | OFF }
  | ONLINE = { ON | OFF }
  | ALLOW_ROW_LOCKS = { ON | OFF }
  | ALLOW_PAGE_LOCKS = { ON | OFF }
  | MAXDOP = max_degree_of_parallelism
}
```

<http://usa.redirectme.net/repriser/data/indexdefrag.html>

Update statistics.



6. Primary Key Clustered Index

```
USE [AdventureWorksDW2008R2]
GO

/***** Object: Table [dbo].[Table_1]    Script Date: 04/27/2010 10:42:16 *****/
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER ON
GO


CREATE TABLE [dbo].[Table_1](
    [NAME] [nchar](10) NOT NULL,
    [JobDescription] [nchar](10) NULL,
    [PositionTitle] [nchar](255) NULL,
    CONSTRAINT [PK_Table_1] PRIMARY KEY CLUSTERED
(
    [NAME] ASC
)WITH (PAD_INDEX = OFF, STATISTICS_NORECOMPUTE = OFF, IGNORE_DUP_KEY = OFF, ALLOW_ROW_LOCKS = ON, ALLOW_PAGE_LOCKS = ON) ON [PRIMARY]
) ON [PRIMARY]
GO
```

Primary Key has to be unique and not null.

Clustered Index is recommended to be on sequential values to avoid fragmentation.

7. Uniqueidentifier

	Column Name	Data Type	Allow Nulls
	NAME	nchar(10)	<input type="checkbox"/>
▶	RID	uniqueidentifier	<input type="checkbox"/>
	JobDescription	nchar(10)	<input checked="" type="checkbox"/>
	PositionTitle	nchar(255)	<input checked="" type="checkbox"/>
			<input type="checkbox"/>

Column Properties	
	
(General)	
(Name)	RID
Allow Nulls	No
Data Type	uniqueidentifier
Default Value or Binding	
Table Designer	
Collation	<database default>
Computed Column Specification	
Condensed Data Type	uniqueidentifier
Description	
Deterministic	Yes
DTS-published	No
Full-text Specification	
Has Non-SQL Server Subscriber	No
Identity Specification	
Indexable	Yes
Is Columnset	No
Is Sparse	No
Merge-published	No
Not For Replication	No
Replicated	No
RowGuid	No
Size	16

`NEWSEQUENTIALID ()`

`NEWSEQUENTIALID()` can only be used with DEFAULT constraints on table columns of type **uniqueidentifier**. For example:

```
CREATE TABLE myTable (ColumnA uniqueidentifier DEFAULT NEWSEQUENTIALID())
```

When `NEWSEQUENTIALID()` is used in DEFAULT expressions, it cannot be combined with other scalar operators. For example, you cannot execute the following:


```
CREATE TABLE myTable (ColumnA uniqueidentifier DEFAULT dbo.myfunction(NEWSEQUENTIALID()))
```

In the previous example, `myfunction()` is a scalar user-defined scalar function that accepts and returns a **uniqueidentifier** value.

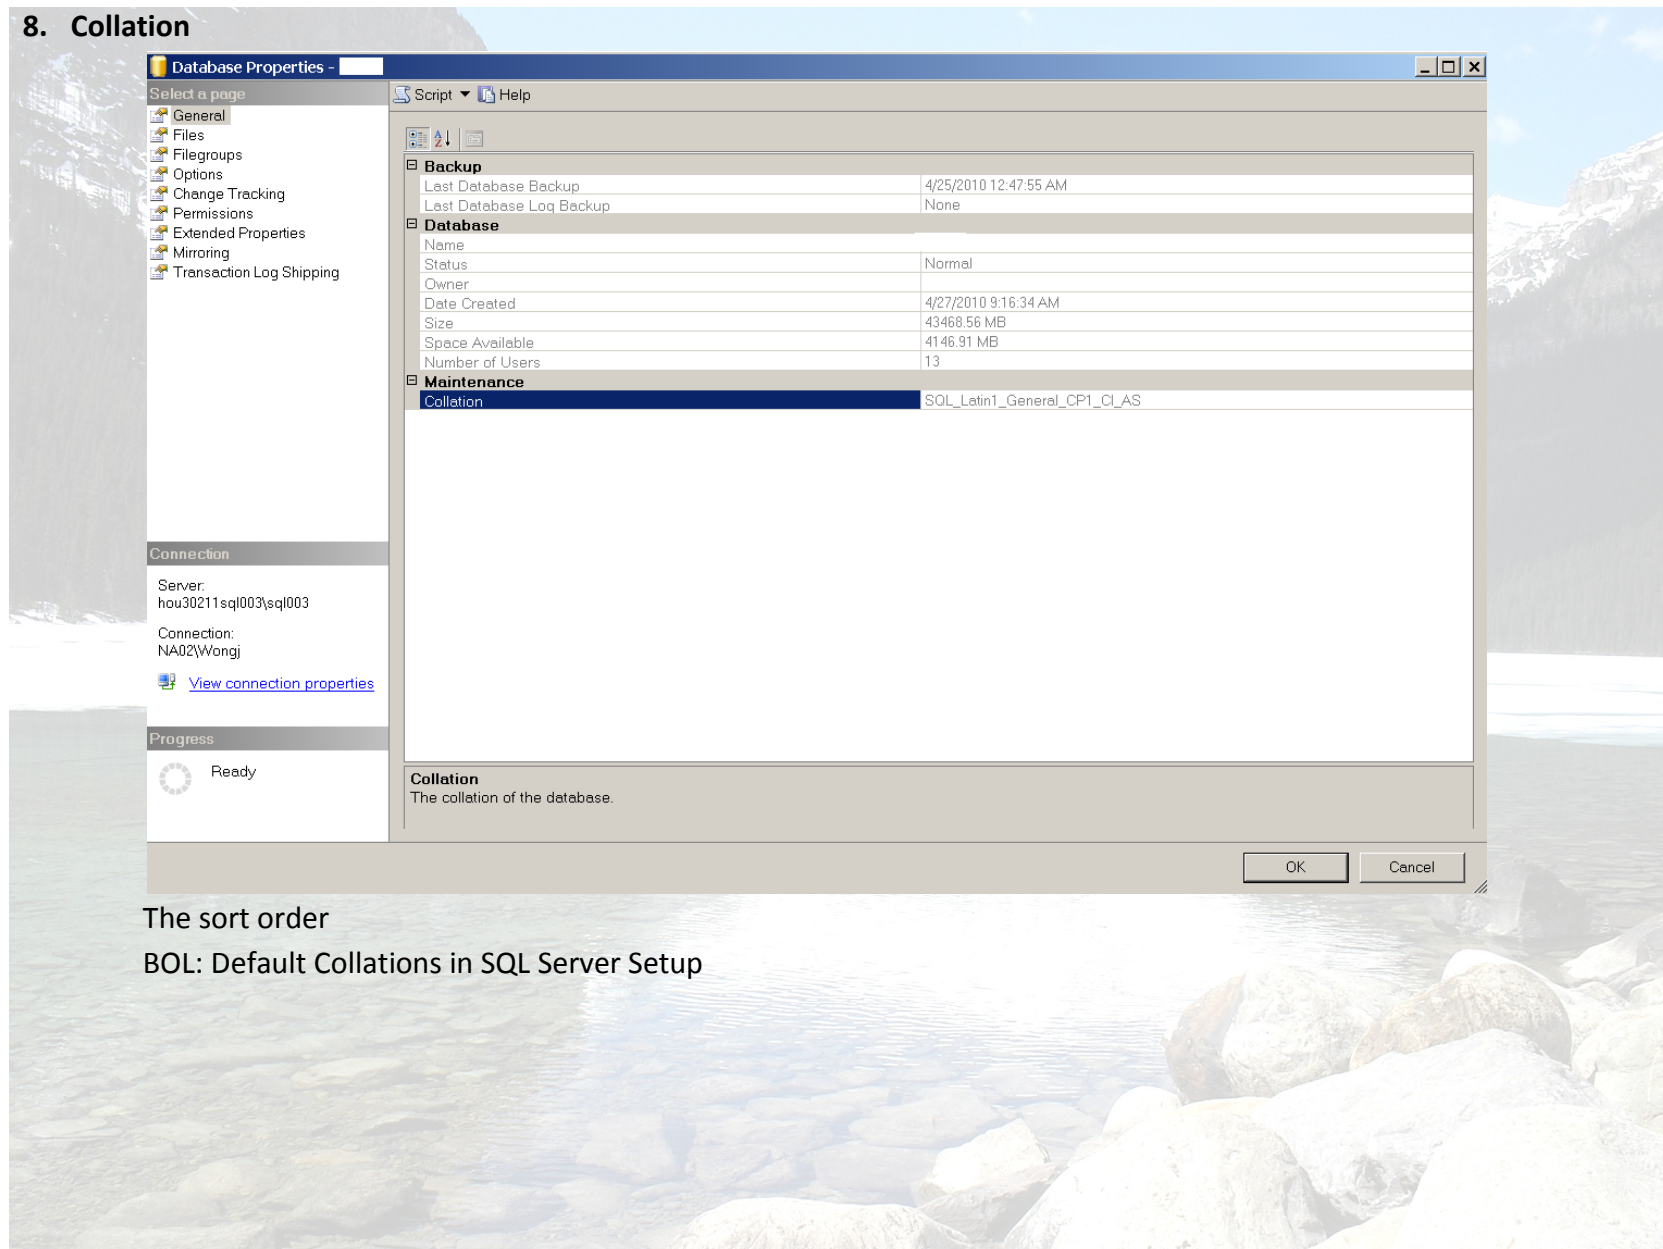
`NEWSEQUENTIALID()` cannot be referenced in queries.

The GUIDs generated by `NEWSEQUENTIALID()` are unique only within a particular computer if the computer does not have a network card.

You can use `NEWSEQUENTIALID()` to generate GUIDs to reduce page contention at the leaf level of indexes.



8. Collation



The sort order

BOL: Default Collations in SQL Server Setup

Dutch (Netherlands)	0x413	Latin1_General_CI_AS
English (Australia)	0xc09	Latin1_General_CI_AS
English (Belize)	0x2809	Latin1_General_CI_AS
English (Canada)	0x1009	Latin1_General_CI_AS
English (Caribbean)	0x2409	Latin1_General_CI_AS
English (Hong Kong S.A.R.)	0x3c09	Latin1_General_CI_AS
English (India)	0x4009	Latin1_General_CI_AS
English (Indonesia)	0x3809	Latin1_General_CI_AS
English (Ireland)	0x1809	Latin1_General_CI_AS
English (Jamaica)	0x2009	Latin1_General_CI_AS
English (Malaysia)	0x4409	Latin1_General_CI_AS
English (New Zealand)	0x1409	Latin1_General_CI_AS
English (Philippines)	0x3409	Latin1_General_CI_AS
English (Singapore)	0x4809	Latin1_General_CI_AS
English (South Africa)	0x1c09	Latin1_General_CI_AS
English (Trinidad)	0x2c09	Latin1_General_CI_AS
English (United Kingdom)	0x809	Latin1_General_CI_AS
English (United States)	0x409	SQL_Latin1_General_CP1_CI_AS
English (Zimbabwe)	0x3009	Latin1_General_CI_AS
Estonian	0x425	Estonian_CI_AS

US English; North America West Europe;

US-Europe-MiddleEast-FarEast-Russia

9. Data type – nvarchar(255), the “N”

Column Name	Data Type	Allow Nulls
NAME	nchar(10)	<input checked="" type="checkbox"/>
	hierarchyid	<input type="checkbox"/>
	image	
	int	
	money	
	nchar(10)	
	ntext	
	numeric(18, 0)	
	nvarchar(50)	



Unicode Character Strings

nchar	ntext
nvarchar	

10. Allow Nulls

Column Name	Data Type	Allow Nulls
NAME	nchar(10)	<input checked="" type="checkbox"/>
	hierarchyid	<input type="checkbox"/>
	image	
	int	
	money	
	nchar(10)	
	ntext	
	numeric(18, 0)	
	nvarchar(50)	



What is NULL? Technical definition:

Do we need NULL in database?

What is the impact?

What is the solution?